Refine Search

Search Results -

Terms	Documents
L3 and (shift$4 near8 redundant$3 near4 data)	1

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search: L4

Refine Search

Recall Text

Clear

Interrupt

Search History

DATE: Tuesday, December 07, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
	DB=USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ		
L4	L3 and (shift$4 near8 redundant$3 near4 data)	1	L4
L3	(storage or memory or disk or DASD or RAID) near8 (redundant$3 near4 (stored or writ$4 or storing or transfer$4))	626	L3
	DB=PGPB,USPT; PLUR=YES; OP=ADJ		
L2	L1 and (shift$4 near8 redundant$3 near4 data)	29	L2
L1	(storage or memory or disk or DASD or RAID) near8 (redundant$3 near4 (stored or writ$4 or storing or transfer$4))	2195	L1

END OF SEARCH HISTORY

Your search matched **0** of **1099723** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or enteri new one in the text box.

(storage or memory or disk or dasd or raid) <near/8> Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

---

Results:
No documents matched your query.

Terms used
storage or **memory** or **disk** or **dasd** or **raid** **near/8** **redundant** **near/4** **stored** or **writ** or **storing** or **transfer** and **s**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanc
Try this searcl

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown

1 RAID: high-performance, reliable secondary storage
Peter M. Chen, Edward K. Lee, Garth A. Gibson, Randy H. Katz, David A. Patterson
June 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 2
Full text available: pdf(3.60 MB) Additional Information: full citation, abstract, references, citings,

Disk arrays were proposed in the 1980s as a way to use parallelism between multiple disks to imp Today they appear in the product lines of most major computer manufacturers. This article gives a arrays and provides a framework in which to organize current and future work. First, the article in reviews the driving forces that have popularized disk arrays: performance and reliability. It discus

Keywords: RAID, disk array, parallel I/O, redundancy, storage, striping

2 Performance evaluation of extended storage architectures for transaction processing
Erhard Rahm
June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international cc data**, Volume 21 Issue 2
Full text available: pdf(1.47 MB) Additional Information: full citation, abstract, references, index te

The use of non-volatile semiconductor memory within an extended storage hierarchy promises sig for transaction processing. Although page-addressable semiconductor memories like extended me caches are commercially available since several years, no detailed investigation of their use for tra performed so far. We present a comprehensive simulation study that compares the performance o

3 EVENODD: an optimal scheme for tolerating double disk failures in RAID architectures
M. Blaum, J. Brady, J. Bruck, J. Menon
April 1994 **ACM SIGARCH Computer Architecture News , Proceedings of the 21ST annual Computer architecture**, Volume 22 Issue 2
Full text available: pdf(893.35 KB) Additional Information: full citation, abstract, references, citings,

We present a novel method, that we call EVENODD, for tolerating up to two disk failures in RAID a known scheme for tolerating double disk failures that is optimal with regard to both storage and p addition of only two redundant disks and consists of simple exclusive-OR computations. A major a only requires parity hardware, which is typically present in standard RAID-5 controllers. Hence, EV

4 Disk cache—miss ratio analysis and design considerations
Alan J. Smith

August 1985 **ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 3

Full text available: pdf(3.13 MB) Additional Information: full citation, abstract, references, citings,

The current trend of computer system technology is toward CPUs with rapidly increasing processin rapidly increasing density, but with disk performance increasing very slowly if at all. The implicatic point the processing power of computer systems will be limited by the throughput of the input/out problem, which is described and evaluated in this paper, is disk cache

5 StorHouse metanoia - new applications for database, storage & data warehousing

Felipe Cariño, Pekka Kostamaa, Art Kaufmann, John Burgess

May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international cc data**, Volume 30 Issue 2

Full text available: pdf(597.88 KB) Additional Information: full citation, abstract, references, index te

This paper describes the StorHouse/Relational Manager (RM) database system that uses and explc active storage hierarchy, we mean that StorHouse/RM executes SQL queries *directly* against data (i.e. disk, optical, and tape) without post processing a file or a DBA having to manage a data set. StorHouse/RM features and internals. We also describe how StorHouse/RM differs from traditional

6 Hierarchical disk cache management in RAID 5 controller

Jung-ho Huh, Tae-mu Chang

December 2003 **The Journal of Computing in Small Colleges**, Volume 19 Issue 2

Full text available: pdf(137.71 KB) Additional Information: full citation, abstract, references, index te

In RAID system, disk cache is one of the important elements in improving the system performanc performance in comparison to single cache and is effective in temporal and spatial locality. The pr two levels. The first level cache is a set associative cache with small block size whereas the secon cache with large block size. In this paper, a RAID 5 disk cache model is presented that is located c

7 Parallelism in relational data base systems: architectural issues and design approaches

Hamid Pirahesh, C. Mohan, Josephine Cheng, T. S. Liu, Pat Selinger

July 1990 **Proceedings of the second international symposium on Databases in parallel**

Full text available: pdf(2.50 MB) Additional Information: full citation, abstract, references, citings,

With current systems, some important complex queries may take days to complete because of: (1 processed, (2) limited aggregate resources. Introducing parallelism addresses the first problem. C resources solve the second problem. According to a survey by Brodie,1 only 10% of computerized argument for both more variety and volume of data to be moved into data base systems. We conj

8 A case for redundant arrays of inexpensive disks (RAID)

David A. Patterson, Garth Gibson, Randy H. Katz

June 1988 **ACM SIGMOD Record , Proceedings of the 1988 ACM SIGMOD international cc data**, Volume 17 Issue 3

Full text available: pdf(1.20 MB) Additional Information: full citation, abstract, references, citings,

Increasing performance of CPUs and memories will be squandered if not matched by a similar perf capacity of Single Large Expensive Disks (SLED) has grown rapidly, the performance improvement Redundant Arrays of Inexpensive Disks (RAID), based on the magnetic disk technology developed attractive alternative to SLED, promising improvements of an order of magnitude in performance,

9 Tolerating multiple failures in RAID architectures with optimal storage and uniform declusteri

Guillermo A. Alvarez, Walter A. Burkhard, Flaviu Cristian

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual Computer architecture**, Volume 25 Issue 2

Full text available: pdf(1.50 MB) Additional Information: full citation, abstract, references, citings,

We present DATUM, a novel method for tolerating multiple disk failures in disk arrays. DATUM is tl

mask any given number of failures, requires an optimal amount of redundant storage space, and s uniformly over disks in the presence of failures without needing large layout tables in controller m information dispersal, a coding technique that admits an efficient hardware implementation. As t .

10 Parity logging overcoming the small write problem in redundant disk arrays
Daniel Stodolsky, Garth Gibson, Mark Holland
May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual Computer architecture**, Volume 21 Issue 2
Full text available: pdf(1.35 MB) Additional Information: full citation, abstract, references, citings,

Parity encoded redundant disk arrays provide highly reliable, cost effective secondary storage with accesses and large write accesses. Their performance on small writes, however, is much worse tha highly reliable, but expensive organization for secondary storage. Unfortunately, small writes are a workload of many important, demanding applications such as on-line transaction processing. This

11 The HP AutoRAID hierarchical storage system
John Wilkes, Richard Golding, Carl Staelin, Tim Sullivan
February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1
Full text available: pdf(1.82 MB) Additional Information: full citation, abstract, references, citings,

Configuring redundant disk arrays is a black art. To configure an array properly, a system adminis of both the array and the workload it will support. Incorrect understanding of either, or changes in to poor performance. We present a solution to this problem: a two-level storage hierarchy implem controller. In the upper level of this hierarchy, two copies of active data are stored to provide f ...

Keywords: RAID, disk array, storage hierarchy

12 Serverless network file systems
Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe, David A. Patterson, Drew S. Roselli, Rand
February 1996 **ACM Transactions on Computer Systems (TOCS)**, Volume 14 Issue 1
Full text available: pdf(2.69 MB) Additional Information: full citation, abstract, references, citings,

We propose a new paradigm for network file system design: serverless network file systems. Whil rely on a central server machine, a serverless system utilizes workstations cooperating as peers tc Any machine in the system can store, cache, or control any block of data. Our approach uses this combination with fast local area networks, to provide better performance and scalability th ...

Keywords: RAID, log cleaning, log structured, log-based striping, logging, redundant data storag

13 The HP AutoRAID hierarchical storage system
J. Wilkes, R. Golding, C. Staelin, T. Sullivan
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM s systems principles**, Volume 29 Issue 5
Full text available: pdf(1.60 MB) Additional Information: full citation, references, citings, index terr

14 Serverless network file systems
T. E. Anderson, M. D. Dahlin, J. M. Neefe, D. A. Patterson, D. S. Roselli, R. Y. Wang
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM s systems principles**, Volume 29 Issue 5
Full text available: pdf(2.48 MB) Additional Information: full citation, references, citings, index terr

15 External memory algorithms and data structures: dealing with massive (

Jeffrey Scott Vitter
June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2

Full text available: pdf(828.46 KB) Additional Information: full citation, abstract, references, citings,

Data sets in large applications are often too massive to fit completely inside the computers interna input/output communication (or I/O) between fast internal memory and slower external memory (performance bottleneck. In this article we survey the state of the art in the design and analysis of algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O cost

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hie access methods, multilevel memory, online, out-of-core, secondary storage, sorting

16 The Zebra striped network file system

John H. Hartman, John K. Ousterhout
August 1995 **ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 3

Full text available: pdf(2.76 MB) Additional Information: full citation, abstract, references, citings,

Zebra is a network file system that increases throughput by striping the file data across multiple s file separately, Zebra forms all the new data from each client into a single stream, which it then st a log-structured file system. This provides high performance for writes of small files as well as for Zebra also writes parity information in each stripe in the style of RAID disk arrays; this ...

Keywords: RAID, log-based striping, log-structured file system, parity computation

17 The Zebra striped network file system

John H. Hartman, John K. Ousterhout
December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM systems principles**, Volume 27 Issue 5

Full text available: pdf(1.93 MB) Additional Information: full citation, abstract, references, citings,

Zebra is a network file system that increases throughput by striping file data across multiple serve separately, Zebra forms all the new data from each client into a single stream, which it then stripe log-structured file system. This provides high performance for writes of small files as well as for re also writes parity information in each stripe in the style of RAID disk arrays; this increase ...

18 Simulation of computer systems and applications

William S. Keezer
December 1997 **Proceedings of the 29th conference on Winter simulation**

Full text available: pdf(742.22 KB) Additional Information: full citation, references, citings, index terr

19 Parity logging disk arrays

Daniel Stodolsky, Mark Holland, William V. Courtright, Garth A. Gibson
August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3

Full text available: pdf(1.98 MB) Additional Information: full citation, abstract, references, citings,

Parity-encoded redundant disk arrays provide highly reliable, cost-effective secondary storage witl large writes. Their performance on small writes, however, is much worse than mirrored disks—the expensive organization for secondary storage. Unfortunately, small writes are a substantial portior

important, demanding applications such as on-line transaction processing. This paper presents
Keywords: RAID, disk arrays

20 Storage: FAB: building distributed enterprise disk arrays from commodity components
Yasushi Saito, Svend Frølund, Alistair Veitch, Arif Merchant, Susan Spence
October 2004 **Proceedings of the 11th international conference on Architectural support fo operating systems**
Full text available: pdf(671.67 KB) Additional Information: full citation, abstract, references, index te

This paper describes the design, implementation, and evaluation of a Federated Array of Bricks (F provides the reliability of traditional enterprise arrays with lower cost and better scalability. FAB is small storage appliances containing commodity disks, CPU, NVRAM, and network interface cards. I based algorithm to replicate or erasure-code logical blocks across bricks and a reconfigurati ...

Keywords: consensus, disk array, erasure coding, replication, storage, voting